

Support-Graph Preconditioners for 2-Dimensional Trusses

Samuel I. Daitch*

Department of Computer Science
Yale University

Daniel A. Spielman†

Department of Computer Science
Yale University

February 1, 2008

Abstract

We use support theory, in particular the fretsaw extensions of Shklarski and Toledo [ST06a], to design preconditioners for the stiffness matrices of 2-dimensional truss structures that are stiffly connected. Provided that all the lengths of the trusses are within constant factors of each other, that the angles at the corners of the triangles are bounded away from 0 and π , and that the elastic moduli and cross-sectional areas of all the truss elements are within constant factors of each other, our preconditioners allow us to solve linear equations in the stiffness matrices to accuracy ϵ in time $O(n^{5/4}(\log^2 n \log \log n)^{3/4} \log(1/\epsilon))$.

1 Preconditioning

When solving a linear system in an $n \times n$ positive semidefinite matrix A , the running time of an iterative solver can often be sped up by **supporting** A with another matrix B , called a **preconditioner**. An effective preconditioner B has the properties that it is much easier to solve than A , and that A has a low condition number relative to B .

We define here generalized eigenvalues and condition numbers:

Definition 1.1. *For positive semidefinite A, B , the **maximum eigenvalue**, **minimum eigenvalue**, and **condition number** of A relative to B are defined respectively as*

$$\begin{aligned}\lambda_{\max}(A, B) &= \max_{\mathbf{x}: \mathbf{x} \perp \text{null}(B)} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T B \mathbf{x}} \\ \lambda_{\min}(A, B) &= \min_{\mathbf{x}: \mathbf{x} \perp \text{null}(A)} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T B \mathbf{x}} \\ \kappa(A, B) &= \lambda_{\max}(A, B) / \lambda_{\min}(A, B)\end{aligned}$$

where $\mathbf{x} \perp \text{null}(S)$ means that \mathbf{x} is orthogonal to the null space of S .

Note that the standard condition number of A can be expressed as $\kappa(A) = \kappa(A, I)$.

The conjugate gradient method is an example of a linear solver that can be sped up using a preconditioner. The precise analysis of the running time can be found, for example, in [Axe85]:

*Partially supported by NSF grant CCR-0324914.

†Partially supported by NSF grant CCR-0324914.

Theorem 1.2 ([Axe85]). *For positive semidefinite A, B , and vector \mathbf{b} , let \mathbf{x} satisfy $A\mathbf{x} = \mathbf{b}$. Each iteration of the preconditioned conjugate gradient method multiplies one vector by A , solves one linear system in B , and performs a constant number of vector additions. For $\epsilon > 0$, it requires at most $O(\sqrt{\kappa(A, B)} \log(1/\epsilon))$ such iterations to produce a $\tilde{\mathbf{x}}$ that satisfies*

$$\|\tilde{\mathbf{x}} - \mathbf{x}\|_A \leq \epsilon \|\mathbf{x}\|_A$$

1.1 Using a Larger Matrix

In certain situations it may be easier to find a good preconditioner for a matrix A if we treat A as being larger than it really is. That is, if we pad A with zeros to form a larger square matrix $A' = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$, it may be simpler to find a good preconditioner B for A' . We then need to show how to use B to yield a preconditioner for the original matrix A . To this end, we define the Schur complement:

Definition 1.3. *For square matrices A and $B = \begin{bmatrix} B_{11} & B_{12} \\ B_{12}^T & B_{22} \end{bmatrix}$, where square submatrix B_{11} is the same size as A , and such that B_{22} is nonsingular, the **Schur complement** of B with respect to A is*

$$B_S = B_{11} - B_{12}B_{22}^{-1}B_{12}^T$$

While B_S will not automatically be a good preconditioner for A simply because B is a good preconditioner for A' , we do know that the maximum eigenvalue will be the same:

Lemma 1.4. *For positive semidefinite A, B , $\lambda_{max}(A', B) = \lambda_{max}(A, B_S)$.*

We also know that solving a linear system in B_S is as easy as solving a linear system in B :

Lemma 1.5. *$B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix}$ implies $B_S \mathbf{x} = \mathbf{b}$*

For completeness, we give proofs for these lemmas in Appendix A.

1.2 Congestion-Dilation

Suppose that we have matrices A and B that can be expressed as the sums of other matrices, i.e. $A = \sum_i A_i$ and $B = \sum_j B_j$, and that we know how to support each A_i by a subset of the B_j matrices. In this situation, we can use the following lemma to show how B supports A :

Lemma 1.6 (Congestion-Dilation Lemma). *Given the symmetric positive semidefinite matrices $A_1, \dots, A_n, B_1, \dots, B_m$ and $A = \sum_i A_i$ and $B = \sum_j B_j$ and given sets $\Sigma_i \subseteq [1, \dots, m]$ and real values s_i that satisfy*

$$\lambda_{max}(A_i, \sum_{j \in \Sigma_i} B_j) \leq s_i$$

it holds that

$$\lambda_{max}(A, B) \leq \max_j \left(\sum_{i:j \in \Sigma_i} s_i \right)$$

A proof of this lemma is given in Appendix A. It is an adaptation of, for example, Proposition 9.4 in [BH03]. The expression $\sum_{i,j \in \Sigma_i} s_i$ can be thought of as the **congestion** of matrix B_j , analogous to the concept of congestion for graph embeddings which we define in Section 4.

2 Trusses and Stiffness Matrices

Definition 2.1. A **2-dimensional truss** $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$ is an undirected weighted planar graph with vertices $[n] = \{1, \dots, n\}$ and edges E , with vertex $i \in [n]$ embedded at point $\mathbf{v}_i \in \mathbb{R}^2$. We allow multiple vertices to be embedded at the same point.

An edge $e = (i, j) \in E$, also called a **truss element**, represents a straight idealized bar from \mathbf{v}_i to \mathbf{v}_j , with positive weight $\gamma(e)$ denoting the product of the bar's cross-sectional area and the elastic modulus of its material. A **truss face** is a triple $\{i, j, k\}$ such that $\{(i, j), (i, k), (j, k)\} \subseteq E$ and no vertex is in the interior of the triangle formed by $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$. Every truss element is required to be contained in some truss face.

There is a particular type of linear system that arises when analyzing the forces on a truss using the finite element method. We define here the type of matrix we wish to solve:

Definition 2.2. Given a truss $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$, for each truss element $e = (i, j) \in E$ we define a length $2n$ column vector $\mathbf{u}_e = [u_e^1 \dots u_e^{2n}]^T$ with 4 nonzero entries satisfying $[u_e^{2i-1} \ u_e^{2i}]^T = -[u_e^{2j-1} \ u_e^{2j}]^T = \frac{\mathbf{v}_i - \mathbf{v}_j}{|\mathbf{v}_i - \mathbf{v}_j|}$, and we define the $2n \times 2n$ matrix

$$A_e = \frac{\gamma(e)}{|\mathbf{v}_i - \mathbf{v}_j|} \mathbf{u}_e \mathbf{u}_e^T$$

The **stiffness matrix** of the truss is then given by:

$$A_{\mathcal{T}} = \sum_{e \in E} A_e$$

Note that a stiffness matrix is positive semidefinite, since for all \mathbf{x} we have

$$\mathbf{x}^T A_{\mathcal{T}} \mathbf{x} = \sum_{e=(\mathbf{v}_i, \mathbf{v}_j) \in E} \frac{\gamma(e)}{|\mathbf{v}_i - \mathbf{v}_j|} \mathbf{x}^T \mathbf{u}_e \mathbf{u}_e^T \mathbf{x} = \sum_{e=(\mathbf{v}_i, \mathbf{v}_j) \in E} \frac{\gamma(e)}{|\mathbf{v}_i - \mathbf{v}_j|} (\mathbf{x}^T \mathbf{u}_e)^2 \geq 0$$

We would like to restrict our attention to trusses with a unique, well-behaved stress-free position. To this end, we make the following definitions:

Definition 2.3. The **rigidity graph** $Q_{\mathcal{T}}$ of a truss \mathcal{T} is the graph with vertex set given by the set of truss faces of \mathcal{T} , and with edges connecting faces that share an edge.

We say that a truss \mathcal{T} is **stiffly-connected** if (1) $Q_{\mathcal{T}}$ is connected, and (2) for every $i \in [n]$, $Q_{\mathcal{T}}^i$ is connected, where $Q_{\mathcal{T}}^i$ is the graph induced by $Q_{\mathcal{T}}$ on the set of faces containing vertex i .

The main contribution of this paper is an algorithm **TrussSolver** for solving linear systems in stiffness matrices of stiffly-connected trusses. We will describe the algorithm later, but we state here the result of our analysis of the running time:

Theorem 2.4 (Main Result). For any stiffly-connected truss $\langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$ such that

- all truss elements have lengths in the range $[l_{min}, l_{max}]$

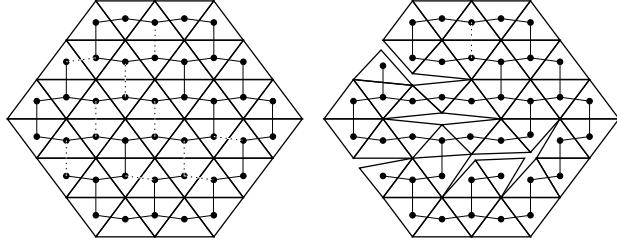


Figure 1: The truss on the right is a fretsaw extension of the truss on the left, as given by the `fretsaw` algorithm. The vertex positions in the fretsaw extension are distorted slightly so as to be able to distinguish vertex copies in the same location. The subgraph F is shown as solid lines, while the rest of the trusses' connectivity graphs are shown as dotted lines. Note that the connectivity graph of the fretsaw extension has one edge not in F .

- all angles of truss faces are in the range $[\theta_{\min}, \pi - \theta_{\min}]$.
- all weights are in the range $[\gamma_{\min}, \gamma_{\max}]$.

for positive constants $l_{\min}, l_{\max}, \theta_{\min}, \gamma_{\min}, \gamma_{\max}$, `TrussSolver` solves linear systems in matrix $A_{\mathcal{T}}$ within relative error ϵ in time $O(n^{5/4}(\log^2 n \log \log n)^{3/4} \log(1/\epsilon))$

2.1 Fretsaw Extension

We will precondition the stiffness matrix using a fretsaw extension, a technique described in [ST06a]. The fretsaw extension of a truss is a new truss created by splitting some of the vertices into multiple copies, without changing the identity of the truss faces.

Definition 2.5. Let $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$ and $\mathcal{T}' = \langle m, \{\mathbf{v}'_i\}_{i=1}^m, E', \gamma' \rangle$ be 2-dimensional trusses, with $m > n$. Let F and F' be the sets of truss faces of \mathcal{T} and \mathcal{T}' respectively. Let $\rho : F' \rightarrow F$ be a bijection between the faces, and let $\pi : [m] \rightarrow [n]$ be a surjection on the vertices. \mathcal{T}' is the (ρ, π) -fretsaw extension of \mathcal{T} if

- for all $i \in [m]$, the i th vertex of \mathcal{T}' is a copy of the $\pi(i)$ th vertex of \mathcal{T} , i.e. $\mathbf{v}'_i = \mathbf{v}_{\pi(i)}$
- for all faces $f = (i, j, k) \in F'$, the vertices of face f are copies of the vertices of $\rho(f)$, i.e. $\rho(i, j, k) = (\pi(i), \pi(j), \pi(k))$
- every edge $(i, j) \in E'$ has the same weight as the edge of which it is a copy, i.e. $\gamma'(i, j) = \gamma(\pi(i), \pi(j))$

Since each vertex in \mathcal{T} has at least one copy in \mathcal{T}' , we follow that convention that $\forall i \in [n] : \pi(i) = i$. Thus, for $i \in [n]$, vertex i in \mathcal{T}' can be considered the “original copy” of vertex i in \mathcal{T} .

A fretsaw extension has the following property which makes it a useful preconditioner:

Lemma 2.6 (see [ST06a] Lemma 8.14). Let A, B be the stiffness matrices of $\mathcal{T}, \mathcal{T}'$ respectively. Let B_S be the Schur complement of B with respect to A .

If \mathcal{T}' is a fretsaw extension of \mathcal{T} , then $\lambda_{\min}(A, B_S) \geq 1/2$

We give a short proof of this lemma in Appendix A. Combining this with Lemma 1.4, we find:

Corollary 2.7. *Let A, B be the stiffness matrices of $\mathcal{T}, \mathcal{T}'$ respectively. Let B_S be the Schur complement of B with respect to A , and let $A' = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$ be the same size as B .*

If \mathcal{T}' is a fretsaw extension of \mathcal{T} , then $\kappa(A, B_S) \leq 2\lambda_{\max}(A', B)$

Now, consider a truss \mathcal{T} and fretsaw extension \mathcal{T}' , with respective rigidity graphs $Q_{\mathcal{T}}$ and $Q_{\mathcal{T}'}$. By construction, every pair of faces that share an edge in \mathcal{T}' must also share an edge in \mathcal{T} . That is, if we let $\rho(Q_{\mathcal{T}'}) = \{(\rho(f_1), \rho(f_2)) : (f_1, f_2) \in Q_{\mathcal{T}'}\}$ denote the graph isomorphic to $Q_{\mathcal{T}'}$ on the faces of \mathcal{T} , then $\rho(Q_{\mathcal{T}'}) \subseteq Q_{\mathcal{T}}$.

As it turns out, for any subgraph $H \subseteq Q_{\mathcal{T}}$ of our choice, we can construct a fretsaw extension with $Q_{\mathcal{T}'}$ (almost) isomorphic to H . We present a linear-time construction here. For technical reasons, this construction also takes as input a map $\tau : [n] \rightarrow F$ that for each vertex in \mathcal{T} specifies one face containing that vertex. The construction ensures that the face in \mathcal{T}' corresponding to $\tau(i)$ contains the original copy of vertex i . This feature will be useful later, and does not diminish the generality of the algorithm.

Lemma 2.8. *There exists a linear-time algorithm $\langle \mathcal{T}', \rho \rangle = \text{fretsaw}(\mathcal{T}, H, \tau)$ that takes a stiffly-connected truss $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$ with face set F , a connected spanning subgraph H of $Q_{\mathcal{T}}$, and a map $\tau : [n] \rightarrow F$ from each vertex to a truss face containing it, and returns a stiffly-connected (ρ, π) -fretsaw extension $\mathcal{T}' = \langle m, \{\mathbf{v}_i\}_{i=1}^m, E', \gamma' \rangle$ satisfying:*

1. for all $i \in [n]$, $i \in \rho^{-1}(\tau(i))$
2. $H' \subseteq Q'$
3. if $|H'| = n - 1 + k$ (i.e. H' is a spanning tree plus k additional edges), then $|Q' - H'| \leq k$

where Q' denotes $Q_{\mathcal{T}'}$. and H' denotes $\rho^{-1}(H) = \{(f_1, f_2) \in Q' : (\rho(f_1), \rho(f_2)) \in H\}$

Proof. Here is the construction, an example of which is given in Figure 1:

fretsaw(\mathcal{T}, H, τ)

First, for each vertex i in \mathcal{T} , we create the set $\pi^{-1}(i)$ of copies of vertex i in \mathcal{T}' : (Recall that we call $i \in \pi^{-1}(i)$ the “original copy”.)

- Let F_i denote the set of faces of \mathcal{T} containing vertex i , and let H_i denote the graph induced by H on F_i . For each connected component of H_i , we put one copy of vertex i in \mathcal{T}' . The original copy is assigned to the connected component of H_i containing face $\tau(i)$.

Now, for $f \in F_i$, let $\phi(i, f) \in \pi^{-1}(i)$ denote the copy of vertex i that is assigned to the component of H_i containing f . It is straightforward to construct the faces of \mathcal{T}' :

- For each face $f = (i, j, k)$ of \mathcal{T} , we create a face $\rho^{-1}(f) = (\phi(i, f), \phi(j, f), \phi(k, f))$ in \mathcal{T}' .

The first property is directly enforced by the construction.

To see why the second property holds, consider an edge $(f_1, f_2) \in H$, where (i, j) is the edge shared by faces f_1 and f_2 . Since edge (f_1, f_2) is present in both H_i and H_j , faces $\rho^{-1}(f_1)$ and $\rho^{-1}(f_2)$ will share the same copies of vertices i and j , and so they too will share an edge.

As for the third property, suppose that H' has $n - 1 + k$ edges, and thus divides the plane into k regions. It suffices to show that each such region contains at most one edge in $Q' - H'$.

Let (f_1, f_2) be an edge in $Q_{\mathcal{T}'} - H'$. Let (i, j) be the edge shared by f_1 and f_2 . Let Q'_i and $H'_i \subseteq Q'_i$ denote the graphs induced on F_i by Q' and H' respectively.

Since f_i and f_j share the vertices i and j , we know there must a path from f_1 to f_2 both in H'_i and in H'_j . Of course neither path contains the edge (f_i, f_2) , since it is not in H' . The only possibility then is that Q'_i is a cycle $H'_i \cup \{(f_1, f_2)\}$, where H'_i is a path from f_1 to f_2 , and similarly for Q'_j . Thus $Q'_i \cup Q'_j = H'_i \cup H'_j \cup \{(f_1, f_2)\}$, and so (f_1, f_2) is the only edge of Q' inside the region enclosed by cycle $H'_i \cup H'_j$. \square

3 Path Lemma

We will need to construct a fretsaw extension with a truss matrix that can be solved quickly. In particular, the fretsaw extension we construct will have a connectivity graph that is close to a spanning tree (i.e. close to having $n - 1$ edges), because we can efficiently find a sparse Cholesky factorization of its truss matrix. The following result is proven in Appendix (?somewhere?):

Lemma 3.1. *Let A be the stiffness matrix of an n -vertex truss \mathcal{T} , where $Q_{\mathcal{T}}$ comprises a spanning tree R plus a set S of additional edges. A Cholesky factorization $A = PLL^TP^T$ can be found in time $O(n + |S|^{3/2})$, where P is a permutation matrix, and such that systems in lower triangular matrix L can be used to solve systems in A in time $O(n + |S| \log |S|)$.*

Now, of course we want to construct a fretsaw extension whose truss matrix provides good support for the original truss matrix. If we can give a supporting subset of faces in the fretsaw extension for each element in the original truss, then we can use Lemma 1.6 to bound the maximum generalized eigenvalue. To this end, we show that a simply-connected set of faces that connects two vertices supports the matrix of an element between the pair of vertices proportionally to the cube of the number of faces:

Lemma 3.2 (Path Lemma). *Let $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$ be a k -face simply-connected truss, and let e_0 be a truss element with weight $\gamma(e_0)$ between any pair $\mathbf{v}_p, \mathbf{v}_q \in V$ of vertices in the truss, such that in \mathcal{T} and e_0*

- all truss elements have lengths in the range $[l_{\min}, l_{\max}]$
- all angles of truss faces are in the range $[\theta_{\min}, \pi - \theta_{\min}]$.
- all weights are in the range $[\gamma_{\min}, \gamma_{\max}]$.

for positive constants $l_{\min}, l_{\max}, \theta_{\min}, \gamma_{\min}, \gamma_{\max}$. Then:

$$\lambda_{\max}(A_{e_0}, A_{\mathcal{T}}) = O(k^3)$$

We first note the simply-connected truss \mathcal{T} must contain a simply-connected subset of faces whose rigidity graph is a path, such that the first face in the path is the only one containing vertex p and the last triangle is the only one containing vertex q . Thus, without loss of generality we may assume that \mathcal{T} is itself such a “truss path”, because removing the extra faces can only increase the value of λ_{\max} .

Let us then number the faces in the path in order f_1, \dots, f_n , and let us number the $n + 2$ vertices as follows:

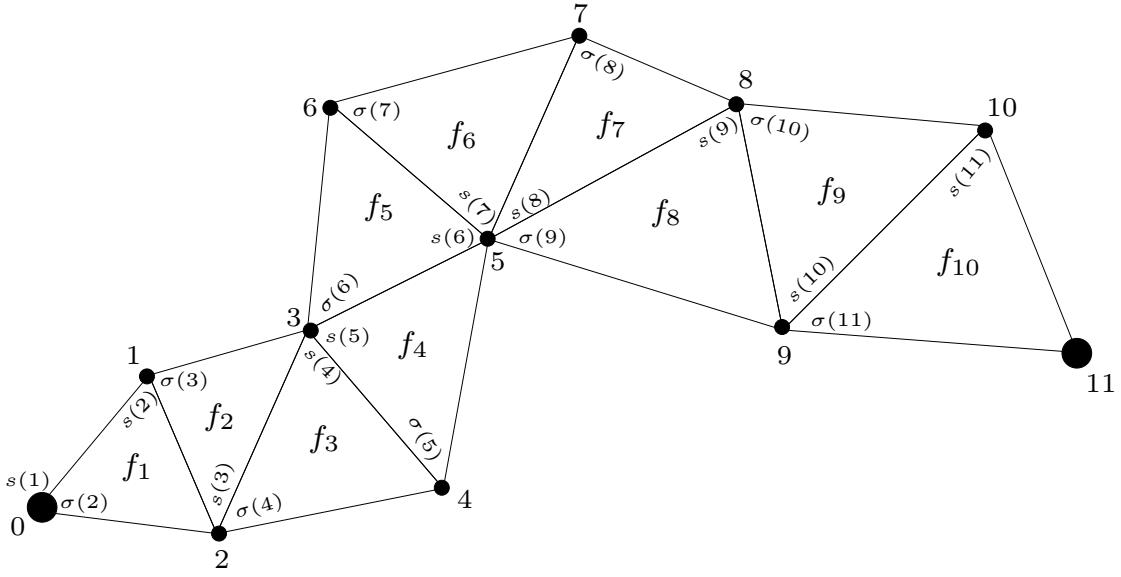


Figure 2: A truss path from v_0 to v_{11} , with triangles and vertices labeled appropriately.

- Let $p = 0$ be the vertex in f_1 but not f_2 .
- Let vertices 1 and 2 be the pair of vertices shared by f_1 and f_2 .
- For $3 \leq i \leq n + 1$, let i be the vertex in f_{i-1} but not f_i . (In particular, $q = n + 1$.)

Furthermore, for $2 \leq i \leq n$, consider the labeling of the three vertices in f_{i-1}

- i is one of the vertices shared by f_{i-1} and f_i .
- Let $s(i)$ denote the other vertex shared by f_{i-1} and f_i .
- Let $\sigma(i)$ denote the vertex in f_{i-1} but not f_i .

For completeness, define $s(1) = 0, s(n+1) = n, \sigma(n+1) = n-1$, so for $1 \leq i \leq n$:

$$f_i = \{s(i), i, i+1\} = \{\sigma(i+1), s(i+1), i+1\}$$

and the set of truss elements is given by:

$$E = \{(0, 1)\} \cup \bigcup_{i=2}^{n+1} \{(s(i), i), (\sigma(i), i)\}$$

An example of this labeling is given in Figure 2.

In the proof we make use of the following canonical definition of a perpendicular vector:

Definition 3.3. *The counterclockwise perpendicular of a vector $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \in \mathbb{R}^2$ is*

$$\mathbf{x}^\perp = \begin{bmatrix} -x_2 \\ x_1 \end{bmatrix}$$

We note some useful properties of the perpendicular:

Claim 3.4. *For any $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^2$:*

1. $(-\mathbf{x})^\perp = -(\mathbf{x}^\perp)$
2. $(\mathbf{x} + \mathbf{y})^\perp = \mathbf{x}^\perp + \mathbf{y}^\perp$
3. $|\mathbf{x}^T \mathbf{y}^\perp| = |\mathbf{x}| |\mathbf{y}| |\sin \theta|$ where θ is the difference in angle between \mathbf{x} and \mathbf{y}
4. $\frac{(\mathbf{x}-\mathbf{y})^T (\mathbf{x}-\mathbf{z})^\perp}{(\mathbf{x}-\mathbf{y})^T (\mathbf{y}-\mathbf{z})^\perp} = 1$
5. $\frac{\mathbf{x}^\perp \mathbf{y}^T}{\mathbf{y}^T \mathbf{x}^\perp} + \frac{\mathbf{y}^\perp \mathbf{x}^T}{\mathbf{x}^T \mathbf{y}^\perp} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I$

Proof. The first three properties are trivial.

The fourth can easily be seen by $\frac{(\mathbf{x}-\mathbf{y})^T (\mathbf{x}-\mathbf{z})^\perp}{(\mathbf{x}-\mathbf{y})^T (\mathbf{y}-\mathbf{z})^\perp} = \frac{(\mathbf{x}-\mathbf{y})^T (\mathbf{x}-\mathbf{y})^\perp}{(\mathbf{x}-\mathbf{y})^T (\mathbf{y}-\mathbf{z})^\perp} + \frac{(\mathbf{x}-\mathbf{y})^T (\mathbf{y}-\mathbf{z})^\perp}{(\mathbf{x}-\mathbf{y})^T (\mathbf{y}-\mathbf{z})^\perp} = 0 + 1$

Here is a proof of the fifth:

$$\frac{\mathbf{x}^\perp \mathbf{y}^T}{\mathbf{y}^T \mathbf{x}^\perp} + \frac{\mathbf{y}^\perp \mathbf{x}^T}{\mathbf{x}^T \mathbf{y}^\perp} = \frac{1}{x_2 y_1 - x_1 y_2} \begin{bmatrix} x_2 y_1 & x_2 y_2 \\ -x_1 y_1 & -x_1 y_2 \end{bmatrix} - \frac{1}{x_2 y_1 - x_1 y_2} \begin{bmatrix} x_1 y_2 & x_2 y_2 \\ -x_1 y_1 & -x_2 y_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I$$

□

Now, let $\mathbf{x}^* = [x_0^*, x_1^*, \dots, x_{2n-1}^*]^T$ be a vector that maximizes $\frac{\mathbf{x}^T A_{e_0} \mathbf{x}}{\mathbf{x}^T A_T \mathbf{x}}$ over all values of \mathbf{x} . Let \mathbf{x}_i^* denote $[x_{2i-1}^*, x_{2i}^*]^T$

In particular we choose an \mathbf{x}^* such that $(\mathbf{x}_1^* - \mathbf{x}_0^*)$ is parallel to $(\mathbf{v}_1 - \mathbf{v}_0)$, by taking advantage of the following property of the null space:

Lemma 3.5. Define $\mathbf{x}^R = [x_0^R, \dots, x_{2n-1}^R]^T$ to be the vector satisfying $[x_{2i-1}^R, x_{2i}^R]^T = (\mathbf{v}_i - \mathbf{v}_0)^\perp$. \mathbf{x}^R is in the null space of both A_{e_0} and A_T .

Proof. For the matrix A_e of any single element $e = (i, j)$, we have:

$$A_e \mathbf{x}^R = \frac{\gamma}{|\mathbf{v}_i - \mathbf{v}_j|} \mathbf{u}_e \mathbf{u}_e^T \mathbf{x}^R = \frac{\gamma}{|\mathbf{v}_i - \mathbf{v}_j|^3} \mathbf{u}_e \left((\mathbf{v}_i - \mathbf{v}_j)^T (\mathbf{v}_i - \mathbf{v}_j)^\perp + (\mathbf{v}_j - \mathbf{v}_i)^T (\mathbf{v}_j - \mathbf{v}_i)^\perp \right) = 0$$

□

Note that we can eliminate the component of \mathbf{x}^* perpendicular to $(\mathbf{v}_1 - \mathbf{v}_0)$ by adding the appropriate multiple of \mathbf{x}^R .

Now, let us focus momentarily on a single vertex i , and the two elements $(s(i), i)$ and $(\sigma(i), i)$ that connect vertex i to lower numbered vertices. The terms $\mathbf{x}^T A_{(s(i), i)} \mathbf{x}$ and $\mathbf{x}^T A_{(\sigma(i), i)} \mathbf{x}$ are zero respectively when

$$(\mathbf{v}_i - \mathbf{v}_{s(i)})^T (\mathbf{x}_i - \mathbf{x}_{s(i)}) = 0$$

and

$$(\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_i - \mathbf{x}_{\sigma(i)}) = 0$$

Supposing we set $\mathbf{x}_{s(i)} = \mathbf{x}_{s(i)}^*$ and $\mathbf{x}_{\sigma(i)} = \mathbf{x}_{\sigma(i)}^*$, we would like to define \mathbf{d}_i to be the vector such that setting $\mathbf{x}_i = \mathbf{x}_i^* - \mathbf{d}_i$ satisfies both of the above equations.

In particular, we define the vectors

$$\mathbf{d}_i = \begin{cases} \mathbf{x}_1^* - \mathbf{x}_0^* & i = 1 \\ \mathbf{x}_i^* - \mathbf{x}_{s(i)}^* + R_i(\mathbf{x}_{s(i)}^* - \mathbf{x}_{\sigma(i)}^*) & 2 \leq i \leq n+1 \end{cases}$$

where

$$R_i = \frac{(\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T}{(\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp}$$

We claim that these satisfy the following properties:

Lemma 3.6. *The following are properties of the \mathbf{d}_i s:*

1. For all $(j, i) \in E$, $j < i$:

$$(\mathbf{v}_i - \mathbf{v}_j)^T (\mathbf{x}_i^* - \mathbf{x}_j^*) = (\mathbf{v}_i - \mathbf{v}_j)^T \mathbf{d}_i$$

2. For all i :

$$|\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*| \leq \frac{l_{max}}{l_{min} \sin \theta_{min}} \sum_{j=1}^i |\mathbf{d}_j|$$

Proof of 1. The statement is trivial for element $(0, 1)$. There are two other types of elements we must consider: $(s(i), i)$ and $(\sigma(i), i)$.

For an element $(s(i), i)$ we have:

$$\begin{aligned} (\mathbf{v}_i - \mathbf{v}_{s(i)})^T \mathbf{d}_i &= (\mathbf{v}_i - \mathbf{v}_{s(i)})^T (\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*) + \frac{(\mathbf{v}_i - \mathbf{v}_{s(i)})^T (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^\perp (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_{s(i)}^* - \mathbf{x}_{\sigma(i)}^*)}{(\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp} \\ &= (\mathbf{v}_i - \mathbf{v}_{s(i)})^T (\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*) + 0 \end{aligned}$$

using the fact that $(\mathbf{v}_i - \mathbf{v}_{s(i)})^T (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp = 0$

For an element $(\sigma(i), i)$ we have:

$$\begin{aligned} (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T \mathbf{d}_i &= (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*) + \frac{(\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_{s(i)}^* - \mathbf{x}_{\sigma(i)}^*)}{(\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp} \\ &= (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*) + (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_{s(i)}^* - \mathbf{x}_{\sigma(i)}^*) \\ &= (\mathbf{v}_i - \mathbf{v}_{\sigma(i)})^T (\mathbf{x}_i^* - \mathbf{x}_{\sigma(i)}^*) \end{aligned}$$

□

Proof of 2. For $i \geq 2$, using the fact that $\{s(i), \sigma(i)\} = \{s(i-1), i-1\}$, we have

$$\mathbf{x}_i^* - \mathbf{x}_{s(i)}^* = \mathbf{d}_i - R_i(\mathbf{x}_{s(i)}^* - \mathbf{v}_{\sigma(i)}^*) = \mathbf{d}_i \pm R_i(\mathbf{x}_{i-1}^* - \mathbf{v}_{s(i-1)}^*)$$

Since $\mathbf{x}_1^* - \mathbf{x}_0^* = \mathbf{d}_1$, we recursively find that

$$|\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*| \leq |\mathbf{d}_i| + \sum_{j=1}^{i-1} |(R_i R_{i-1} \cdots R_{j+1}) \mathbf{d}_j|$$

The following finishes the proof:

$$\begin{aligned}
& |R_i R_{i-1} \cdots R_{j+1} \mathbf{d}_j| \\
&= \left| (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp \prod_{k=j+2}^i \left(\frac{(\mathbf{v}_k - \mathbf{v}_{\sigma(k)})^T (\mathbf{v}_{k-1} - \mathbf{v}_{s(k-1)})^\perp}{(\mathbf{v}_k - \mathbf{v}_{\sigma(k)})^T (\mathbf{v}_k - \mathbf{v}_{s(k)})^\perp} \right) \cdot \frac{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T \mathbf{d}_j}{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T (\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)})^\perp} \right| \\
&= \left| (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp \prod_{k=j+2}^i \left(\frac{(\mathbf{v}_k - \mathbf{v}_{\sigma(k)})^T (\mathbf{v}_{s(k)} - \mathbf{v}_{\sigma(k)})^\perp}{(\mathbf{v}_k - \mathbf{v}_{\sigma(k)})^T (\mathbf{v}_k - \mathbf{v}_{s(k)})^\perp} \right) \cdot \frac{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T \mathbf{d}_j}{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T (\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)})^\perp} \right|
\end{aligned}$$

by the fact that $\{s(i), \sigma(i)\} = \{s(i-1), i-1\}$

$$= \left| (\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp \prod_{k=j+2}^i (\pm 1) \cdot \frac{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T \mathbf{d}_j}{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T (\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)})^\perp} \right|$$

by Claim 3.4.4

$$\begin{aligned}
&= \left| \frac{(\mathbf{v}_i - \mathbf{v}_{s(i)})^\perp (\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T \mathbf{d}_j}{(\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)})^T (\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)})^\perp} \right| \\
&\leq \frac{|\mathbf{v}_i - \mathbf{v}_{s(i)}| |\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)}| |\mathbf{d}_j|}{|\mathbf{v}_{j+1} - \mathbf{v}_{\sigma(j+1)}| |\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)}| \sin \theta_{min}}
\end{aligned}$$

by Claim 3.4.3

$$\begin{aligned}
&= \frac{|\mathbf{v}_i - \mathbf{v}_{s(i)}| |\mathbf{d}_j|}{|\mathbf{v}_{j+1} - \mathbf{v}_{s(j+1)}| \sin \theta_{min}} \\
&\leq \frac{l_{max}}{l_{min} \sin \theta_{min}} |\mathbf{d}_j|
\end{aligned}$$

□

To finish proving the path lemma, we will need to use the following fact:

Lemma 3.7. *Let $\mathbf{u}_1, \mathbf{u}_2$ be unit vectors whose angles differ by θ . Then for any \mathbf{v} and $a, b > 0$:*

$$\frac{(\mathbf{u}_1^T \mathbf{v})^2}{a} + \frac{(\mathbf{u}_2^T \mathbf{v})^2}{b} \geq \frac{\sin^2 \theta |\mathbf{v}|^2}{a+b}$$

Proof. Let α be the angle between \mathbf{u}_1 and \mathbf{v} .

We must show that $\frac{1}{a} \cos^2 \alpha + \frac{1}{b} \cos^2(\alpha + \theta) \geq \frac{1}{a+b} \sin^2 \theta$ □

Recall that we wish to prove

$$\lambda_{max}(A_{e_0}, A_{\mathcal{T}}) = \frac{(\mathbf{x}^*)^T A_{e_0} \mathbf{x}^*}{(\mathbf{x}^*)^T A_{\mathcal{T}} \mathbf{x}^*} = O(n^3)$$

Let us first bound the denominator $(\mathbf{x}^*)^T A_{\mathcal{T}} \mathbf{x}^*$.

$$\begin{aligned}
& (\mathbf{x}^*)^T A_{\mathcal{T}} \mathbf{x}^* \\
&= (\mathbf{x}^*)^T A_{(0,1)} \mathbf{x}^* + \sum_{i=2}^{n+1} ((\mathbf{x}^*)^T A_{(s(i),i)} \mathbf{x}^* + (\mathbf{x}^*)^T A_{(\sigma(i),i)} \mathbf{x}^*) \\
&\geq \gamma_{min} \left[\frac{\left[\left(\frac{\mathbf{v}_1 - \mathbf{v}_0}{|\mathbf{v}_1 - \mathbf{v}_0|} \right)^T (\mathbf{x}_1^* - \mathbf{x}_0^*) \right]^2}{|\mathbf{v}_1 - \mathbf{v}_0|} + \sum_{i=2}^{n+1} \left(\frac{\left[\left(\frac{\mathbf{v}_i - \mathbf{v}_{s(i)}}{|\mathbf{v}_i - \mathbf{v}_{s(i)}|} \right)^T (\mathbf{x}_i^* - \mathbf{x}_{s(i)}^*) \right]^2}{|\mathbf{v}_i - \mathbf{v}_{s(i)}|} + \frac{\left[\left(\frac{\mathbf{v}_i - \mathbf{v}_{\sigma(i)}}{|\mathbf{v}_i - \mathbf{v}_{\sigma(i)}|} \right)^T (\mathbf{x}_i^* - \mathbf{x}_{\sigma(i)}^*) \right]^2}{|\mathbf{v}_i - \mathbf{v}_{\sigma(i)}|} \right) \right] \\
&= \gamma_{min} \left[\frac{|\mathbf{d}_1|^2}{|\mathbf{v}_1 - \mathbf{v}_0|} + \sum_{i=2}^{n+1} \left(\frac{\left[\left(\frac{\mathbf{v}_i - \mathbf{v}_{s(i)}}{|\mathbf{v}_i - \mathbf{v}_{s(i)}|} \right)^T \mathbf{d}_i \right]^2}{|\mathbf{v}_i - \mathbf{v}_{s(i)}|} + \frac{\left[\left(\frac{\mathbf{v}_i - \mathbf{v}_{\sigma(i)}}{|\mathbf{v}_i - \mathbf{v}_{\sigma(i)}|} \right)^T \mathbf{d}_i \right]^2}{|\mathbf{v}_i - \mathbf{v}_{\sigma(i)}|} \right) \right]
\end{aligned}$$

by Lemma 3.6 and the fact that $\mathbf{d}_1 = \mathbf{x}_1^* - \mathbf{x}_0^*$ is parallel to $(\mathbf{v}_1 - \mathbf{v}_0)$

$$\geq \gamma_{min} \left[\frac{|\mathbf{d}_1|^2}{|\mathbf{v}_1 - \mathbf{v}_0|} + \sum_{i=2}^{n+1} \left(\frac{\sin^2 \theta_{min} |\mathbf{d}_i|^2}{|\mathbf{v}_i - \mathbf{v}_{s(i)}| + |\mathbf{v}_i - \mathbf{v}_{\sigma(i)}|} \right) \right]$$

by Lemma 3.7

$$\geq \sum_{i=1}^{n+1} \left(\frac{\gamma_{min} \sin^2 \theta_{min} |\mathbf{d}_i|^2}{2l_{max}} \right)$$

Next we bound the numerator:

$$\begin{aligned}
(\mathbf{x}^*)^T A_{e_0} \mathbf{x}^* &= \frac{\gamma(e_0)}{|\mathbf{v}_{n+1} - \mathbf{v}_0|} \left[\left(\frac{\mathbf{v}_{n+1} - \mathbf{v}_0}{|\mathbf{v}_{n+1} - \mathbf{v}_0|} \right)^T (\mathbf{x}_{n+1}^* - \mathbf{x}_0^*) \right]^2 \\
&\leq \frac{\gamma_{max}}{l_{min}} |\mathbf{x}_{n+1}^* - \mathbf{x}_0^*|^2 \\
&\leq \frac{\gamma_{max}}{l_{min}} \left(|\mathbf{x}_{n+1}^* - \mathbf{x}_{s(n+1)}^*| + |\mathbf{x}_{s(n+1)}^* - \mathbf{x}_{s(s(n+1))}^*| + \dots + |\mathbf{x}_1^* - \mathbf{x}_0^*| \right)^2 \\
&\leq \frac{\gamma_{max}}{l_{min}} \left((n+1) \frac{l_{max}}{l_{min} \sin \theta_{min}} \sum_{i=1}^{n+1} |\mathbf{d}_i| \right)^2
\end{aligned}$$

by Lemma 3.6

Combining the above, we get:

$$\begin{aligned}
\lambda_{max}(A_{e_0}, A_{\mathcal{T}}) &= \frac{(\mathbf{x}^*)^T A_{e_0} \mathbf{x}^*}{(\mathbf{x}^*)^T A_{\mathcal{T}} \mathbf{x}^*} \\
&\leq \frac{2}{\sin^3 \theta_{min}} \left(\frac{l_{min}}{l_{max}} \right)^3 \frac{\gamma_{max}}{\gamma_{min}} \frac{\left(\sum_{i=1}^{n+1} |\mathbf{d}_i| \right)^2}{\sum_{i=1}^{n+1} |\mathbf{d}_i|^2} (n+1)^2 \\
&\leq \frac{2}{\sin^3 \theta_{min}} \left(\frac{l_{min}}{l_{max}} \right)^3 \frac{\gamma_{max}}{\gamma_{min}} (n+1)^3
\end{aligned}$$

where the last inequality follows from Cauchy-Schwarz.

4 Graph Embeddings

Our remaining task is to describe how to map edges to supporting face sets with low congestion, as required for Lemma 1.6. We need some graph theoretic notions which will inform how we choose to support the edges.

For a graph with edges E , we use $\mathcal{P}(E)$ to denote the power set of E . Thus $\mathcal{P}(E)$ includes all paths in the graph.

Let us define the notion of embedding vertex pairs of a graph into paths in a subgraph:

Definition 4.1. For an unweighted graph (V, E) , a subgraph $H \subseteq E$, and a set Z of pairs of vertices in V , an **embedding** of Z onto H is a map $\pi : Z \rightarrow \mathcal{P}(H)$, where $\pi(v, w) \subseteq H$ is a path in H whose endpoints are v, w .

The **stretch** of π is $str(\pi) = \sum_{z \in Z} |\pi(z)|$.

The **congestion** of π is $cong(\pi) = \max_{f \in H} \left[\sum_{z \in Z: f \in \pi(z)} |\pi(z)| \right]$

A particular example is embedding the edges of a graph onto a spanning tree:

Definition 4.2. For an unweighted connected graph (V, E) and spanning tree $T \subseteq E$, let $T(v, w)$ denote the path in T that connects v to w .

The **stretch** of T is $str(T) = \sum_{e \in E} |T(e)|$.

We will make use of algorithms that take a graph and generate a spanning tree, augmented with a few additional edges, such that given vertex pairs have a low congestion embedding into the augmented tree. First, we need to create a low stretch-spanning tree. The best known result is from [EEST06]:

Theorem 4.3. There exists an algorithm $T = \text{LowStretch}(G)$, that takes a connected graph $G = (V, E)$, runs in time $O(|E| \log^2 |E|)$, and outputs a spanning tree T with stretch $O(|E| \log^2 |V| \log \log |V|)$.

We can then use the low-stretch spanning tree to create an augmented spanning tree with the desired low congestion embedding. This algorithm is given in Appendix B:

Theorem 4.4. There exists an algorithm $S = \text{LowCongestAugment}(G, T, Z, \psi, k)$ that takes a planar graph $G = (V, E)$, a spanning tree T of G , a set Z of pairs of vertices in V , an embedding $\psi : Z \rightarrow \mathcal{P}(E)$, and an integer k . The algorithm runs in time $O(|E| \log |V| + cong(\psi)|E|)$ and returns a set of edges $S \subseteq E$ of size at most k , such that there exists an embedding $\pi : Z \rightarrow \mathcal{P}(T \cup S)$ with congestion $O(\frac{1}{k} str(T) cong(\psi))$.

TrussSolver(\mathcal{T}, k)

Let A be the stiffness matrix of $\mathcal{T} = \langle n, \{v_i\}_{i=1}^n, E, \gamma \rangle$, and let F be its face set.

1. Define $\tau : [n] \rightarrow F$ to map vertex i to an arbitrary face in the set F_i of faces containing i .
 2. Run $R = \text{LowStretch}(Q_{\mathcal{T}})$ in time $O(n \log^2 n)$.
 3. Define $Z = \{(\tau(i), \tau(j)) : (i, j) \in E\}$.
Define an embedding $\psi : Z \rightarrow \mathcal{P}(Q_{\mathcal{T}})$, by defining $\psi(\tau(i), \tau(j))$ to be an arbitrary path from $\tau(i)$ to $\tau(j)$ in $F_i \cup F_j$. (We know F_i and F_j intersect because some face contains edge (i, j) .)
Run $S = \text{LowCongestAugment}(Q_{\mathcal{T}}, R, Z, \psi, k)$ in time $O(n \log n)$.
 4. Run $\langle \mathcal{T}', \rho \rangle = \langle \langle V', E', \gamma' \rangle, \rho \rangle = \text{fretsaw}(\mathcal{T}, R \cup S, \tau)$. Let B be the truss matrix of \mathcal{T}' .
 5. Use Lemma 3.1 to find a Cholesky factorization $B = PLL^T P^T$ in time $O(n + k^{3/2})$, such that L can be used to solve equations in B in time $O(n + k \log k)$.
 6. Run preconditioned conjugate gradient using B_S , the Schur complement of B with respect to A , as the preconditioner. Use L to solve equations in B_S , by solving equations in B (see Lemma 1.5).
- The relative error will be down to ϵ after $O(\sqrt{\kappa(A, B_S)} \log \frac{1}{\epsilon})$ iterations.

Figure 3: The **TrussSolver** algorithm

5 Solving the Linear System

In Figure 3 we present the complete **TrussSolver** algorithm for solving linear systems in matrix A that is the stiffness matrix of truss $\mathcal{T} = \langle n, \{v_i\}_{i=1}^n, E, \gamma \rangle$ with face set F . The algorithm preconditions A using the stiffness matrix B of a fretsaw extension $\mathcal{T}' = \langle m, \{v'_i\}_{i=1}^m, E', \gamma' \rangle$ with face set F' . The algorithm uses a parameter k that will be chosen later. We will show that (with the right choice of k) the algorithm attains a relative error of ϵ in time $O(n^{5/4} (\log^2 n \log \log n)^{3/4} \log \frac{1}{\epsilon})$.

We want to use the Congestion Dilation Lemma (Lemma 1.6) to give an upper bound on $\kappa(A, B_S)$. Recall that for $A' = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$ of the same size as B , Corollary 2.7 says that $\kappa(A, B_S) \leq 2\lambda_{max}(A', B)$. To bound $\lambda_{max}(A', B)$, for each edge $(p, q) \in E$ we will give a face subset $F'_{p,q} \subseteq F'$ connecting p to q , such that this embedding of edges to truss paths has low congestion.

In particular, let $E'_{p,q} \subseteq E'$ be the set of truss elements in the faces $F_{p,q}$, so that $\mathcal{T}'_{p,q} = \langle m, \{v'_i\}, E'_{p,q}, \gamma' \rangle$ is the “subtruss” of \mathcal{T}' comprising the faces $F_{p,q}$. Lemma 3.2 states that

$$\lambda_{max}(A_{(p,q)}, A_{\mathcal{T}'_{p,q}}) = O(|F_{p,q}|^3)$$

so Lemma 1.6 yields

$$\lambda_{max}(A', B) \leq \max_{e \in E'} \sum_{(p,q) \in E : e \in E'_{p,q}} O(|F'_{p,q}|^3)$$

It remains for us to describe the truss paths $\mathcal{T}'_{p,q}$ that yield the desired bound.

Recall that we have constructed a subgraph $R \cup S \in Q_{\mathcal{T}}$, for which the **LowCongestAugment** algorithm guarantees that there exists an embedding $\pi : Z \rightarrow \mathcal{P}(R \cup S)$ of low congestion. Let us denote $\pi_{p,q} = \pi(\tau(p), \tau(q))$, the path in $R \cup S \subseteq Q_{\mathcal{T}}$ connecting $\tau(p)$ to $\tau(q)$. Map this path back into $Q_{\mathcal{T}'}$ to get the path $\pi'_{p,q} = \rho^{-1}(\pi_{p,q}) = \{(\rho^{-1}(f_1), \rho^{-1}(f_2)) : (f_1, f_2) \in \pi_{p,q}\}$. We then form $\mathcal{T}'_{p,q}$ from the set $F'_{p,q}$ of faces in $\pi'_{p,q}$.

Let us first determine the congestion of π more precisely. The algorithms **LowCongestAugment** and **LowStretch** guarantee respectively that $\text{cong}(\pi) = O(\frac{1}{k} \text{str}(R) \text{cong}(\psi))$ and $\text{str}(R) = O(n \log^2 n \log \log n)$.

As for $\text{cong}(\psi)$, we have

$$\text{cong}(\psi) = \max_{q \in Q_{\mathcal{T}}} \sum_{z \in Z: q \in \psi(z)} |\psi(z)| \leq \left(\max_{z \in Z} |\psi(z)| \right) \left(\max_{q \in Q} |\{z \in Z : q \in \psi(z)\}| \right)$$

Now note that for all i , $|F_i| \leq \frac{2\pi}{\theta_{\min}} = O(1)$. Since any $\psi(\tau_i, \tau_j)$ only contains triangles in $T_i \cup T_j$, $|\psi(\tau(i), \tau(j))| = O(1)$, and so $\max_{z \in Z} |\psi(z)| = O(1)$.

Similarly, say that $q \in Q_{\mathcal{T}}$ is a pair of faces sharing the edge (i, j) . Since i and j are the only vertices that the pair of faces have in common, q can only be in a path $\psi(\tau_\alpha, \tau_\beta)$ if one of α or β is i or j . So $|\{z : q \in \psi(z)\}| \leq |\{(\alpha, \beta) \in E : \alpha = i \text{ or } \alpha = j\}| \leq \frac{4\pi}{\theta_{\min}} = O(1)$.

Thus, $\text{cong}(\psi) = O(1)$ and $\text{cong}(\pi) = O(\frac{n}{k} \log^2 n \log \log n)$.

We now have:

$$\begin{aligned} \kappa(A, B_S) &\leq \max_{e \in E'} \sum_{(p,q) \in E: e \in E'_{p,q}} O(|F'_{p,q}|^3) \\ &= O \left(\max_{e \in E'} \sum_{\substack{(p,q) \in E: \\ e \in E'_{p,q}}} |F'_{p,q}| \right)^3 \\ &= O \left(\max_{(f_i, f_j) \in Q_{\mathcal{T}'}} \sum_{\substack{(p,q) \in E: \\ (f_i, f_j) \in \pi'_{p,q}}} |\pi'_{p,q}| \right)^3 \\ &= O \left(\max_{(f_i, f_j) \in R \cup S} \sum_{\substack{(p,q) \in E: \\ (f_i, f_j) \in \pi_{p,q}}} |\pi_{p,q}| \right)^3 \\ &= O \left(\max_{(f_i, f_j) \in R \cup S} \sum_{\substack{z \in Z: \\ (f_i, f_j) \in \pi(z)}} |\pi(z)| \right)^3 \\ &= O(\text{cong}(\pi))^3 \\ &= O \left(\frac{n^3}{k^3} (\log^2 n \log \log n)^3 \right) \end{aligned}$$

Steps 1-4 take time $O(n \log^2 n + k^{3/2})$. Each conjugate gradient iteration takes time $O(n + k \log k)$, and the number of iterations is

$$\sqrt{\kappa(A, B_S)} \log \frac{1}{\epsilon} = O\left(\frac{n^{3/2}}{k^{3/2}} (\log^2 n \log \log n)^{3/2} \right) \log \frac{1}{\epsilon}$$

Thus, our total running time is:

$$O\left(n \log^2 n + k^{3/2} + n^{3/2} k^{-3/2} (\log^2 n \log \log n)^{3/2} (n + k \log k) \log \frac{1}{\epsilon}\right)$$

For $k = n^{5/6} (\log^2 n \log \log n)^{1/2}$ this gives a running time of $O\left(n^{5/4} (\log^2 n \log \log n)^{3/4} \log \frac{1}{\epsilon}\right)$.

References

- [Axe85] O. Axelsson. A survey of preconditioned iterative methods for linear systems of algebraic equations. *BIT Numerical Mathematics*, 25(1):165–187, March 1985.
- [BH03] Erik G. Boman and Bruce Hendrickson. Support theory for preconditioning. *SIAM Journal on Matrix Analysis and Applications*, 25(3):694–717, 2003.
- [EEST06] Michael Elkin, Yuval Emej, Daniel A. Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 2006. To appear.
- [ST06a] Gil Shklarski and Sivan Toledo. Rigidity in finite-element matrices: Sufficient conditions for the rigidity of structures and substructures, 2006.
- [ST06b] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. Available at <http://www.arxiv.org/abs/cs.NA/0607105>, 2006.

A Preconditioning Lemmas

We first prove several lemmas dealing with the Schur complement. Recall that the Schur complement of $B = \begin{bmatrix} B_{11} & B_{12} \\ B_{12}^T & B_{22} \end{bmatrix}$ is $B_S = B_{11} - B_{12}B_{22}^{-1}B_{12}^T$

Lemma A.1. *If B is positive semidefinite then for any \mathbf{x}*

$$\min_{\mathbf{y}} \left([\mathbf{x}^T \quad \mathbf{y}^T] B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right) = \mathbf{x}^T B_S \mathbf{x}$$

Proof.

$$\begin{aligned} [\mathbf{x}^T \quad \mathbf{y}^T] B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} &= \mathbf{x}^T B_{11} \mathbf{x} + \mathbf{x}^T B_{12} \mathbf{y} + \mathbf{y}^T B_{12}^T \mathbf{x} + \mathbf{y}^T B_{22} \mathbf{y} \\ &= \mathbf{x}^T (B_{11} - B_{12}B_{22}^{-1}B_{12}^T) \mathbf{x} + \mathbf{x}^T B_{12}B_{22}^{-1}B_{12}^T \mathbf{x} + \mathbf{x}^T B_{12} \mathbf{y} + \mathbf{y}^T B_{12}^T \mathbf{x} + \mathbf{y}^T B_{22} \mathbf{y} \\ &= \mathbf{x}^T B_S \mathbf{x} + (\mathbf{y} + B_{22}^{-1}B_{12}^T \mathbf{x})^T B_{22} (\mathbf{y} + B_{22}^{-1}B_{12}^T \mathbf{x}) \\ &= \mathbf{x}^T B_S \mathbf{x} + \left[\mathbf{y} + B_{22}^{-1}B_{12}^T \mathbf{x} \right]^T B \begin{bmatrix} 0 \\ \mathbf{y} + B_{22}^{-1}B_{12}^T \mathbf{x} \end{bmatrix} \\ &\geq \mathbf{x}^T B_S \mathbf{x} \end{aligned}$$

The last inequality holds because B is positive semidefinite, and it is an equality for $\mathbf{y} = -B_{22}^{-1}B_{12}^T\mathbf{x}$. \square

Lemma 1.4. *For positive semidefinite A, B , $\lambda_{max}(A', B) = \lambda_{max}(A, B_S)$.*

Proof.

$$\begin{aligned}\lambda_{max}(A', B) &= \max_{\mathbf{x}, \mathbf{y}} \frac{[\mathbf{x}^T \quad \mathbf{y}^T] A' \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}}{[\mathbf{x}^T \quad \mathbf{y}^T] B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}} \\ &= \max_{\mathbf{x}, \mathbf{y}} \frac{\mathbf{x}^T A \mathbf{x}}{[\mathbf{x}^T \quad \mathbf{y}^T] B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}} \\ &= \max_{\mathbf{x}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T B_S \mathbf{x}} \\ &= \lambda_{max}(A, B_S)\end{aligned}$$

The third equality uses Lemma A.1. \square

Lemma 1.5. $B \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix}$ implies $B_S \mathbf{x} = \mathbf{b}$

Proof. Multiplying by $\begin{bmatrix} I & -B_{12}B_{22}^{-1} \\ 0 & I \end{bmatrix}$ gives:

$$\begin{aligned}\begin{bmatrix} I & -B_{12}B_{22}^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ B_{12}^T & B_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} &= \begin{bmatrix} I & -B_{12}B_{22}^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix} \\ \begin{bmatrix} B_S & 0 \\ B_{12}^T & B_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} &= \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix}\end{aligned}$$

\square

Lemma 1.6 (Congestion-Dilation Lemma). *Given the symmetric positive semidefinite matrices $A_1, \dots, A_n, B_1, \dots, B_m$ and $A = \sum_i A_i$ and $B = \sum_j B_j$ and given sets $\Sigma_i \subseteq [1, \dots, m]$ and real values s_i that satisfy*

$$\lambda_{max}(A_i, \sum_{j \in \Sigma_i} B_j) \leq s_i$$

it holds that

$$\lambda_{max}(A, B) \leq \max_j \left(\sum_{i: j \in \Sigma_i} s_i \right)$$

Proof. Let us write $A \preceq B$ to mean that $B - A$ is positive semidefinite. We are given that

$$A_i \preceq s_i \sum_{j \in \Sigma_i} B_j$$

So we have:

$$\begin{aligned} A &= \sum_i A_i \\ &\preceq \sum_i s_i \sum_{j \in \Sigma_i} B_j \\ &= \sum_j B_j \sum_{i:j \in \Sigma_i} s_i \\ &\preceq \max_j \left(\sum_{i:j \in \Sigma_i} s_i \right) \sum_j B_j \\ &= \max_j \left(\sum_{i:j \in \Sigma_i} s_i \right) B \end{aligned}$$

□

Lemma 2.6 (see [ST06a] Lemma 8.14). *Let A, B be the stiffness matrices of $\mathcal{T}, \mathcal{T}'$ respectively. Let B_S be the Schur complement of B with respect to A .*

If \mathcal{T}' is a fretsaw extension of \mathcal{T} , then $\lambda_{\min}(A, B_S) \geq 1/2$

Proof. Suppose that $\mathcal{T} = \langle m, \{\mathbf{v}_i\}_{i=1}^m, E', \gamma' \rangle$ is the (ρ, π) -fretsaw extension of $\mathcal{T} = \langle n, \{\mathbf{v}_i\}_{i=1}^n, E, \gamma \rangle$.

Define M to be the $2n \times 2m$ matrix that for all $i \in [n], j \in [m]$ satisfies

$$\begin{bmatrix} M_{2i-1} & M_{2i} \\ M_{2j-1} & M_{2j} \end{bmatrix} = \begin{cases} I & i = \pi(j) \\ 0 & \text{otherwise} \end{cases}$$

and note that for an element $(i, j) \in E'$, $A_{(\pi(i), \pi(j))} = MA_{(i, j)}M^T$. However any element in \mathcal{T} is part of at most two faces, and so can have at most two copies in \mathcal{T}' . Thus, $A \preceq MBM^T \preceq 2A$.

Recalling that $\pi(i) = i$, we note also that M takes the form $[I \quad M_1]$ for some $(2n-2m) \times 2m$ matrix M_1 . Thus, for any \mathbf{x} , we have:

$$\begin{aligned} 2\mathbf{x}^T A \mathbf{x} &\geq \mathbf{x}^T MBM^T \mathbf{x} \\ &= [\mathbf{x}^T \quad M_1 \mathbf{x}^T] B \begin{bmatrix} \mathbf{x} \\ M_1^T \mathbf{x} \end{bmatrix} \\ &\geq \mathbf{x}^T B_S \mathbf{x} \end{aligned}$$

where the last inequality holds by Lemma A.1. □

B Augmented Spanning Tree

We give a proof of Lemma 4.4, which we restate here for convenience. It is a generalization of an algorithm from [ST06b].

Lemma 4.4. *There exists an algorithm $S = \text{LowCongestAugment}(G, T, Z, \psi, k)$ that takes a planar graph $G = (V, E)$, a spanning tree T of G , a set Z of pairs of vertices in V , an embedding $\psi : Z \rightarrow \mathcal{P}(E)$, and an integer k . The algorithm runs in time $O(|E| \log |V| + \text{cong}(\psi)|E|)$ and returns a set of edges $S \subseteq E$ of size at most k , such that there exists an embedding $\pi : Z \rightarrow \mathcal{P}(T \cup S)$ with congestion $O(\frac{1}{k} \text{str}(T) \text{cong}(\psi))$.*

We need to make use of the following tree decomposition algorithm from [ST06b]:

Theorem B.1 ([ST06b] Theorem 8.3). *There exists a linear-time algorithm*

$$((W_1, \dots, W_c), \rho) = \text{decompose}(T, E, \eta, k)$$

that on input

- a tree T on vertices V
- a set E of edges forming a planar graph on V
- a function $\eta : E \rightarrow \mathbb{R}^+$
- a positive integer $k \leq \sum_{e \in E} \eta(e)$

outputs sets $W_1, \dots, W_c \subseteq V$, where $c \leq k$, and a function ρ that maps each edge in E to either a set or pair of sets in $\{W_1, \dots, W_c\}$ such that:

- $V = \bigcup_{i=1}^c W_i$, and for all $i \neq j$, $|W_i \cup W_j| \leq 1$
- for all i , the graph induced by T on W_i (which we denote T_i) is connected
- for each $(u, v) \in E$, there are i, j (possibly equal) such that $u \in W_i$ and $v \in W_j$ and $\rho(u, v) = \{W_i, W_j\}$
- the graph $\langle [c], \{(i, j) : \exists e \in E \text{ s.t. } \rho(e) = \{W_i, W_j\}\} \rangle$ is planar
- for all W_i such that $|W_i| > 1$,

$$\sum_{e \in E : W_i \in \rho(e)} \eta(e) \leq \frac{4}{k} \sum_{e \in E} \eta(e)$$

We can now prove the lemma.

Proof. Here is the `LowCongestAugment` algorithm:

$S = \text{LowCongestAugment}(\langle V, E \rangle, T, Z, \psi, k)$

1. For each $e \in E$, define $\eta(e) = \sum_{z:e \in \psi(z)} \sum_{e' \in \psi(z)} |T(e')|$
2. Set $((W_1, \dots, W_c), \rho) = \text{decompose}(T, E, \eta, \lfloor \frac{k}{3} \rfloor)$.
3. Let $E_{i,j} = \{e \in E : \rho(e) = \{W_i, W_j\}\}$. For all nonempty $E_{i,j}$ define

$$s_{i,j} = \arg \min_{e \in E_{i,j}} |T(e)|$$

and let S be the set of all the $s_{i,j}$.

Let us analyze the running time. In step 1, we must compute $\eta(e)$ for all $e \in E$. First we compute and record $|T(e)|$ for all $e \in E$. [ST06b] gives a method to do this in time $O(|E| \log |V|)$. We can then compute each $\eta(e)$ by summing $\sum_{z:e \in \psi(z)} |\psi(z)| \leq \text{cong}(\psi)$ of the $|T(e)|$ values. This gives a total time of $O(|E| \log |V| + \text{cong}(\psi)|E|)$ for step 1, and the remaining steps clearly run faster than $O(|E| \log |V|)$.

Let also note that $|S| \leq k$. This follows from the fact that the graph $\langle [c], \{(i, j) : E_{i,j} \neq \emptyset\} \rangle$ is planar and has $|S|$ edges. Since a planar graph on $c \leq k/3$ vertices cannot have more than $3c - 6 < k$ edges, we have that $|S| \leq k$.

Now let us demonstrate the existence of a low congestion embedding. For each $(v, w) \in E$, let us define a path $\pi(v, w)$ in $T \cup S$ from v to w , as follows:

- If $\rho(v, w)$ is a singleton $\{W_i\}$, then we simply define $\pi(v, w) = T(v, w)$

Note that $\pi(v, w) \subseteq T_i$.

- If $\rho(v, w)$ is a pair $\{W_i, W_j\}$, then let $(v', w') = s_{i,j} \in S$ and define $\pi(v, w) = T(v, v') \cup \{(v', w')\} \cup T(w', w)$.

Note that

$$\pi(v, w) \subseteq T_i \cup T_j \cup \{s_{i,j}\}$$

and that

$$|\pi(v, w)| \leq |T(v, w)| + |T(v', w')| + 1 \leq 3|T(v, w)|$$

Furthermore for $(v, w) \in S$, define $\pi(v, w) = \cup_{e \in \psi(v, w)} \pi(e)$.

Fix an $e_0 \in T \cup S$. By construction of π , e can be only in a path $\pi(z)$ if it is either in S or in some subtree T_k .

With this in mind, define i_0 such that T_{i_0} is the subtree containing e . (There is at most one such tree, but if there is none then choose i_0 arbitrarily.) Define i_1, i_2 such that if $e \in S$ then $e = z_{i_1, i_2}$. (If $e \notin S$ then choose i_1, i_2 arbitrarily). Then for any $z \in Z$, if $e \in \pi(z)$ then at least one of the following must hold:

- For some $e \in \psi(z)$, $W_{i_0} \in \rho(e)$.
- For some $e \in \psi(z)$, $\rho(e) = \{W_{i_1}, W_{i_2}\}$. (In particular, $W_{i_1} \in \rho(e)$.)

Thus we can bound the congestion of π on e_0 :

$$\sum_{z:e_0 \in \pi(z)} |\pi(z)| \leq \sum_{z:\{e \in \psi(z) : W_{i_0} \in \rho(e)\} \neq \emptyset} |\pi(z)| \sum_{z:\{e \in \psi(z) : W_{i_1} \in \rho(e)\} \neq \emptyset} |\pi(z)|$$

We note that for any i :

$$\begin{aligned}
\sum_{z: \{e \in \psi(z) : W_i \in \rho(e)\} \neq \emptyset} |\pi(z)| &\leq \sum_{e: W_i \in \rho(e)} \sum_{z: e \in \psi(z)} |\pi(z)| \\
&\leq \sum_{e: W_i \in \rho(e)} \sum_{z: e \in \psi(z)} \sum_{e' \in \psi(z)} |\pi(e')| \\
&\leq 3 \sum_{e: W_i \in \rho(e)} \sum_{z: e \in \psi(z)} \sum_{e' \in \psi(z)} |T(e')| \\
&= 3 \sum_{e: W_i \in \rho(e)} \eta(e) \\
&\leq \frac{12}{k} \sum_{e \in E} \eta(e)
\end{aligned}$$

Thus:

$$\begin{aligned}
\sum_{z: e_0 \in \pi(z)} |\pi(z)| &\leq \frac{24}{k} \sum_{e \in E} \eta(e) \\
&\leq \frac{24}{k} \sum_{e \in E} \sum_{z: e \in \psi(z)} \sum_{e' \in \psi(z)} |T(e')| \\
&= \frac{24}{k} \sum_{e' \in E} |T(e')| \sum_{z: e' \in \psi(z)} |\psi(z)| \\
&\leq \frac{24}{k} \sum_{e' \in E} |T(e')| cong(\psi) \\
&= \frac{24}{k} str(T) cong(\psi)
\end{aligned}$$

□

Since the choice of e_0 was arbitrary, we have

$$cong(\pi) \leq \frac{24}{k} str(T) cong(\psi)$$